

Forthcoming in Theoria. Please cite published version. Some minor corrections made Aug 26.

Strategic content:

Representations of epistemic modality in biosemantics (and success semantics)

Gunnar Björnsson, Stockholm University

Abstract: A central idea in Ruth Millikan's biosemantics is that a representation's content is restricted to conditions required for the normal success of actions that it has as its function to guide. This paper raises and responds to a problem for this idea. The problem is that the success requirement seems to block us from saying that epistemic modal judgments represent our epistemic circumstances. For the normal success of actions guided by these judgments seems to depend on what is actually the case, not on whether or to what extent various possibilities were supported by the evidence. In response, I argue, first, that actions guided by epistemic modal judgments have as their function to implement strategies for handling epistemic circumstances, second, that the successful performance of this function requires that aspects of these circumstances obtain, and, third, that biosemantics can thus understand epistemic modal judgments as representing these aspects. The recognition of such strategic contents introduces complications; I further argue that these are benign.

1. Overview

Ruth Millikan's biosemantic program famously improves on many other forms of naturalistic semantics in two significant ways: by understanding representational content in terms of what is required for the success of systems that are guided by the representations and by understanding the relevant notion of success in terms of etiological functions. In this paper, I first raise what looks like a problem for this program and related forms of success semantics. When we are guided by

epistemic modal judgments, such as judgments about what might be the case or what is likely given the evidence, we seem to be concerned with adjusting our behavior to our epistemic circumstances. One would think, then, that our understanding of these circumstances would serve as a representation of them. But the success of the actions that we perform guided by this understanding does not seem to depend on these circumstances. Based on this, it is tempting to conclude that Millikan's biosemantic program fails to account for what in many ways looks like a clear case of natural representation. In response, I suggest that we should recognize that actions can have as their function to *implement strategies* for handling epistemic circumstances. Since the successful performance of this function depends on aspects of these circumstances, biosemantics allows that epistemic modal judgments represent such aspects.

I proceed as follows: In sections 2, 3, and 4, I briefly introduce the relevant features of biosemantics, raise the problem of epistemic modals, and discuss and reject some initial responses. In section 5, I explain how the recognition of "implementation functions" lets biosemantics provide representational contents for epistemic modal judgments and claims. In section 6, finally, I respond to worries having to do with the fact that epistemic modal judgments and ordinary unguarded beliefs and assertions now all seem to get two contents, as they guide actions that have both implementation functions and ordinary concrete functions.

2. Biosemantics and SUCCESS

The problem that concerns me has its ground in the following implication of Millikan's biosemantics:

SUCCESS: The representational content of some item is a condition for the normal success of further items (states, behaviors) that the item has as its function to guide.

Though the discussion will focus primarily on Millikan's biosemantic understanding of SUCCESS, I take the central problem and the general form of the solution that I will propose to be relevant for

the wider class of “success semantics”.¹ Moreover, problems for SUCCESS are of general interest, as the condition has considerable pre-theoretic appeal. For ordinary beliefs about our surroundings, for example, it seems very plausible that the normal success of the actions or omissions they guide rely on their truth. My belief that there is a sharp kitchen knife on the dining room table might lead me to walk there to get it to prepare food or to remove it to protect children playing in the room, or prevent me from looking for it elsewhere. Under normal circumstances, attempts at these actions to facilitate food preparation, prevent children from being hurt, or avoid wasting energy are successful only insofar as the knife is actually on the kitchen table.

Biosemantics promises to make systematic sense of SUCCESS in relation to beliefs. Moreover, it does so in a way that extends in a principled fashion to a wide range of phenomena, allowing us to highlight important similarities as well as bring interesting differences and relations more clearly into view. For illustration, consider one of Millikan’s recurring examples: the beaver’s splash (see e.g. Millikan 1989: 288). It seems to warn nearby beavers of predators, prompting them to hide underwater. The hiding behavior successfully protects a beaver from danger in normal ways only

¹ For Millikan’s view, see e.g. Millikan 1984; 1989; 2004; 2017. Philosophers who have later tried to develop forms of success semantics often say that it originated with a brief passage from Frank Ramsey (1927: 159), though Ramsey calls it “the pragmatist view”. For non-biosemantic defenses, see e.g. Whyte 1990; Blackburn 2005; Nanay 2013. For another biosemantic defense, see Papineau 1993: 69–77. For criticism, see Brandom 1994; Tang 2014. Ramsey and Nanay both formulate their version of success semantics for simpler representational systems than explicit human beliefs; Whyte, Papineau, and Blackburn are primarily concerned with the latter. By contrast, Millikan’s account is completely general.

Though there are other potential problems with Millikan’s biosemantic program, I will here assume that they can be handled without significant modification, or at least in ways that leave intact the core aspects of biosemantics and SUCCESS that are relevant to my discussion. For an overview of issues concerning the relevant etiological notion of function in relation to the semantic program, see Häggqvist 2013. For a recent collection of critical essays and responses from Millikan, see Ryder et al. 2012.

insofar as there are predators in the vicinity. According to SUCCESS, this makes the presence of predators a candidate representational content of the splash. Or consider ordinary descriptive utterance in the declarative. Arguably, such utterances have as their function to produce beliefs in hearers, beliefs that will serve their function of guiding action in normal ways only when they are true. The general biosemantic story thus straightforwardly connects the representational content of the utterance to that of the corresponding beliefs (see e.g. Millikan 1984: 54ff).

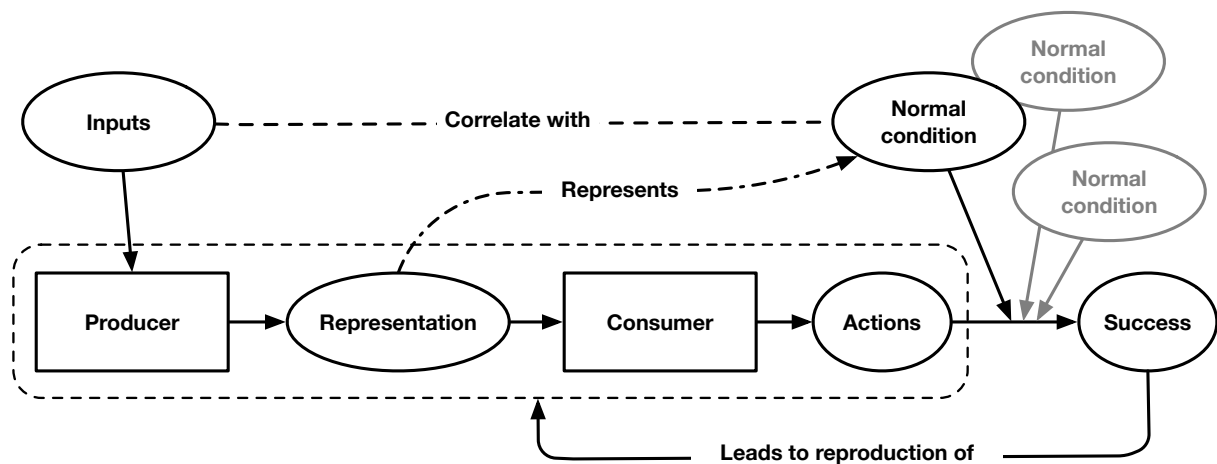
As Millikan has forcefully argued, a biosemantic version of SUCCESS has a solution to what otherwise looks like an insurmountable problem for naturalistic theories of meaning. Such theories have sometimes tried to explain how states or events can be about or represent something in terms of whether they are caused by, or carry natural information about, or are counterfactually dependent on what is represented. The problem with these attempts is two-fold. On the one hand, most causal, informational, or counterfactual relations, even those produced by biological systems, are not plausibly representational. Strong restrictions need to be in place to avoid overgenerating representations. On the other hand, representations are in general not only fallible but oftentimes very unreliable in getting things right, and often correlate much more reliably with other states than those represented, in particular their various proximate causes. The relevant restrictions thus need to allow that representational relations are causally, informationally, or counterfactually weak.

Biosemantics deals with this problem in two ways. First, it understands representations as things that are produced by one system (the “producer”) for the use by another system (the “consumer”) to guide “actions” of the latter because (i) consumer actions normally need for their success for certain states-of-affairs to obtain and (ii) because producer systems have tended (however weakly) to produce items that coincide with these states-of-affairs in virtue of being sensitive to conditions that correlate with the relevant conditions for success. By thus focusing on what the producer needs to provide for normal consumer success—by relying on SUCCESS—biosemantics has the tools to weed out the jungle of non-representational correlations and dependencies.

Second, biosemantics provides a systematic way of understanding references to “function” and “conditions for normal success” in SUCCESS. It understands attributions of functions in etiological

terms: very roughly, functions are effects that ancestors of the current systems or items they produce have had such that these effects explain the reproduction of the systems or items. And it understands conditions for normal success, or “Normal conditions”, in terms of conditions that are part of a minimal general explanation of how something has performed its function on conditions when it has successfully done so (Millikan 1984; 1989; 1993).² (For a simplified schematic overview of biosemantics, see Figure 1.)

Figure 1. Biosemantics, schematic overview



Representational producer-consumer systems are reproductively established because the production of representations is sensitive to inputs that correlate systematically with Normal conditions for the actions that the representations guide. But representations typically represent neither such inputs nor the correlations in question; they only do so when these things have been part of how the relevant actions have normally performed their function. For the beaver’s act of diving underwater to successfully avoid predators, for example, it does not matter what prompted the beaver splash or whether the prompt is correlated with the presence of predators; it only matters whether predators are actually present on the occasion. This is part of why SUCCESS can strongly

² For a slightly more elaborate account of the relevant etiological functions, see n. 5.

reduce the number of viable candidates for representational content. It is also what makes for the problem of epistemic modals.

3. SUCCESS and the problem of epistemic modals

The problem that I will formulate and try to solve is that SUCCESS seems to block plausible accounts of content for epistemic modal judgments. These include judgments that p must happen, or might happen, or is likely to happen, when such judgments seem to track whether p is ensured by, compatible with, or well supported by the evidence available for and against p . The problem is that, as with the beaver splash, these epistemic circumstances do not seem to matter for the success of typical actions undertaken in response to these judgments.

Consider first judgments that p *must (epistemically) be the case*. To apply SUCCESS, we first need a sense of what behavior is guided by such judgments. The most obvious answer is *proceeding as if p* : engaging in practical and theoretical reasoning involving the premise that p . We next need to identify the conditions of normal success for this behavior. The most obvious such condition is simply that p . At a first glance, then, it seems that biosemantics, or success semantics more generally, tells us that insofar as the judgment that it must be that p represents, it represents what the judgment that p represents, namely that p . But at least at a first glance, it seems that whether it must be that p is a matter of how p stands in relation to the available evidence for and against p . This is what we are trying to get right when judging whether p must be the case.

Consider next the judgment that p *might be the case*. It characteristically governs actions designed to pay off when p is the case. Thinking that it might be rainy I bring an umbrella; that pays off in normal ways if it does indeed rain. Thinking that it might be sunny, I bring sun screen; that pays off

if it is sunny. On SUCCESS, the judgment that it might be that p thus seems to represent that p .³ But this seems to misidentify the relevant content. In judging that it might be that p , we empathically fall short of judging that p is the case. Instead we seem to be concerned with whether p is compatible with the available evidence, or something of that sort.

One might think that these worries can be dismissed. There is at least a sense in which someone who thinks that it might be that p and acts accordingly turns out to have made a mistake if p isn't the case. It is also tempting to think of the judgment that it might be that p as simply a *partial belief* that p , in which case one might think that it *should* have the same representational content as a full belief that p .⁴ In defending versions of success semantics, both Whyte (1990: 156–7) and Papineau (1993: 74) thus suggest that the content of partial belief is derived from that of full belief, raising no additional issues (though see Tang 2014 for objections). The suggestion is of course entirely in line with the familiar claim that different degrees of beliefs involve differences in attitude towards one and the same content. There is also good reason to think that partial beliefs cannot be ordinary beliefs about probabilities, that is, beliefs that one can hold to various degrees. (If you believe to degree .3 that the probability of p is around .5 and believe to degree .7 that the probability is higher than .8, then your degree of belief in p is unlikely to be around .5 nor higher than .8.) But one can accept these last claims while insisting that the degree to which one believes that p itself represents a relation between p and one's epistemic circumstances. For, as biosemantics can help us see, not all representations are beliefs, or subject to the same cognitive operations as ordinary beliefs. Moreover, the problem that I am pressing in this section concerns representation generally, not

³ This assumes that judgments that it might be that p have been systematically positively correlated (however weakly) with the truth of p . If we take the relevant contrast to be judgments that *it cannot be that p* , such a correlation should be uncontroversial.

⁴ This suggestion might seem to absurdly imply that the mundane conjunctive judgment that *p might be the case and might not be the case* represents that p and $\text{not-}p$. But that would only follow if the normal success of conjunctive judgments required that the world satisfies the content of each conjunct.

ordinary belief. The problem is that while our epistemic modal judgments seem to have as their function to adjust behaviors to our epistemic circumstances in a way that makes it natural to think of them as representing these circumstances, SUCCESS seems to block us from doing so.

This problem becomes even more vivid when we turn to the judgments that are most plausibly identified with degrees of belief, namely epistemic modal judgment to the effect that *p* is likely to a certain degree. To see what SUCCESS has to say about these, we start by asking what sort of actions such judgments control. Most saliently, judgments attributing likelihoods to *p* determine the weight given in theoretical or practical reasoning to conclusions based on the assumption that *p*, in proportion to how likely *p* is taken to be. Based on such weights, one might then hedge one's bets, look for more information, or simply take a risk. Consider, for example, a dehydrated gazelle. Its actions take into account the strength of evidence that lions are present by the nearby waterhole as well as the value of possible outcomes. (We might think of the values as the weight given to the outcomes when ranking options, and perhaps think of these weights as representing the expected effect on rate of reproduction.) In light of these considerations, the gazelle can go straight for the water, thus minimizing dehydration, or stay away, thus avoiding the effort that would be needed to outrun any lions by the waterhole. It can also take more or less guarded approaches, hedging its bets by making its way towards the water but doing so slowly or cautiously, being ready to turn away should there be lions around. If it does and there are no lions around, it gets its much-needed rehydration, but later than on the unguarded approach. If there are lions around, it survives but exerts more precious energy running for its life than if it had just stayed away.

To illustrate the problem, let us assume that its decision procedure can be modelled as an attempt to maximize expected value, as illustrated by Table 1. The table lists values of outcomes depending on what actions are taken and what the circumstances are, as well as (rounded) expected values given four distributions of likelihoods. (No claims of realism.)

Table 1. Values of outcome and (rounded) expected values given four likelihood distributions

| Approach | Value | | Expected value given likelihood of lions | | | |
|--------------|-------|----------|--|-----------|----------|----------|
| | Lions | No lions | P=.95 | P=.75 | P=.25 | P=.05 |
| Unguarded | -20 | 10 | -19 | -13 | 3 | 9 |
| Guarded | -2 | 8 | -2 | .5 | 6 | 8 |
| Very guarded | -1 | 6 | -.7 | .8 | 4 | 6 |
| None | 0 | 0 | 0 | 0 | 0 | 0 |

Depending on its distribution of likelihoods, the gazelle will approach without caution, opt for the (somewhat) guarded approach, opt for the very guarded approach, or not approach at all. It seems very natural to assume that the choice between strategies is not only causally affected by the epistemic circumstances, but also sensibly adjusts the gazelle’s behavior to fit them. Moreover, it seems very natural to understand the way in which this fit is achieved as involving representations of these circumstances. However, the sort of fit that SUCCESS encodes does not seem to adequately capture what is going on in cases like these. The problem is perhaps the clearest if we compare the two cases of hedging: the guarded and very guarded approaches. First, they seem to have the very same success conditions. Both are partially successful if there are no lions around, as each approach will prevent further dehydration (though less quickly than on the unguarded approach). They are also partially successful if there *are* lions around, as they let the gazelle avoid being killed by the lions (though in a costlier way than by staying away entirely). Second, in neither case is this partial success explained by the antecedent likelihood of lions: what determines whether a certain approach leads to a certain outcome is the actual presence or absence of lions. Nonetheless, the choice between these approaches would seem to be an attempt to adapt to aspects of the circumstances, namely the epistemic circumstances.

To summarize: The problem that epistemic modal judgments pose for biosemantics and SUCCESS stems from

EPISTEMICS LOST: The epistemic circumstances that we attend to when making epistemic modal judgments and seem to adapt to when acting based on such judgments are not Normal conditions for the success of the actions in question.

Together, SUCCESS and EPISTEMICS LOST imply that our judgments about what might, must, or is likely to be the case cannot represent the very aspects of our epistemic circumstances to which our epistemic modal judgments seem to adapt our actions.

4. Some unsatisfactory responses

A first response to the problem stemming from EPISTEMICS LOST would be to acknowledge a certain counterintuitiveness but insist that the Normal conditions discussed above *do* constitute the modal facts in question: judgments that p must be, or that p might be, or that there is some positive likelihood that p all simply represent that p . In favor of this response, it should be acknowledged that counterintuitiveness regarding one category of cases should not be seen as decisive: perhaps SUCCESS and biosemantics have enough independent justification to motivate biting some bullets, especially given how unreliable and conflicted intuitions about modal reality are. More serious is that this response abandons the idea that representations are things by which goal-directed systems adapt to their circumstances because these things tend to be correlated with relevant aspects of those circumstances. I take it that part of the appeal of SUCCESS and biosemantics has been exactly their capacity to make more precise this idea. Other things being equal, a naturalist account of representational content should stay true to this idea also in the case of epistemic modals.

A second response grants that the Normal conditions discussed above cannot provide a range of facts for our various modal judgments to correspond to. Instead, it insists that we can nevertheless make sense of our practice of making epistemic modal judgments in terms of the function that these judgments serve in our wider cognitive economy. Millikan (this issue) has recently taken this route: though she denies that token epistemic modal claims have truth-conditions, “the general practice of using rough statistics based on past experience to set confidence levels pays off over time, leaving us

better prepared for the future, on average, than we would have been otherwise” (10). Notably, Millikan’s claim here is not primarily about epistemic modal *judgments*, but about utterances in the indicative containing epistemic modal expressions; what we might call “modal claims”. Her ground for denying that such claims have truth-conditions is also different than what I’ve outlined above. What she stresses is that although the degrees of confidence corresponding to different degrees of probabilities are based on sensitivity to various relative frequencies, there is no privileged reference class for such frequencies for a given epistemic modal claim that could make the claim correct or incorrect. In spite of these differences, her point is well taken and relevant here: to deny that modal judgments represent modal facts is not to deny that the practice is sensible, or that there is some explanation for its proliferation.

Still, there is something deeply unsatisfactory about the conclusion that when we engage in deliberation under uncertainty, we are not in fact adjusting our actions to our epistemic circumstances based on representations of these. First, I take it to be undeniable that we are in fact trying to adjust to our epistemic circumstances. Indeed, Millikan’s own claim about the practice seems to imply that we are. For suppose, as I think she wants to say, that

STRATEGIC SUCCESS: There is something about the general strategy of being guided by levels of confidence that are sensitive to statistics (such as relative frequencies) that explain why the strategy has been reproduced, and why the strategy has won out over relevant competing strategies.

Then it seems that there has to be some story, however imprecise, about the statistical properties to which confidence levels have corresponded such that this made the strategy superior. Given that there is such a story, it seems to follow that

SYSTEMATIC STRATEGIC FUNCTION: The system that adjusts levels of confidence to the relevant statistical properties and adjusts behaviors to levels of confidence has as its function to do just this: to adjust behaviors to better fit the relevant statistics.

But if we are indeed trying to adjust to our epistemic circumstances when we rely on our epistemic modal judgments, it seems that a naturalistic account of representation should explain how these

judgments represent the relevant aspects of these circumstances. If SUCCESS prevented us from doing this, as our discussion above seemed to suggest, that would be reason to look for other ways of suitably constraining content. Given how important SUCCESS has been in dealing with problems encountered by other forms of naturalistic semantics, this would be bad news.

5. Epistemic circumstances as Normal conditions for strategy implementation

Luckily, biosemantics and SUCCESS *can* account for the fact that epistemic modal judgments represent our epistemic circumstances. The key is provided by SYSTEMATIC STRATEGIC FUNCTION. Given that the system has as its function to implement the strategy, it has as its function to produce actions that relate appropriately to the aspects of the epistemic circumstances that are constituted by the relevant statistics. But then it follows straightforwardly that

IMPLEMENTATION FUNCTION: The actions produced by the system have as their function to implement epistemic strategies: to relate appropriately to the relevant statistics.⁵

⁵ Granted that *the system* has been reproduced for implementing a certain strategy, does it really follow that any *individual implementation* of that strategy has contributed to its reproduction in virtue of being implementations of it? It seems that when the gazelle's cautiously approaching the waterhole contributes to its survival and thus to reproduction of the system, it does so because it results in rehydration or the avoidance of lions, not because it implements some epistemic strategy. (This is just a version of the problem identified in section 3.) Given this, how can the latter be the function of individual implementations? I think that we should set this worry aside. The relevant kind of function here is what Millikan calls a "derived" function, the sort of function satisfying condition (2) in this (condensed) definition of the relevant etiological functions:

... for an item *A* to have a function *F* as a "proper function", it is necessary (and close to sufficient) that one of these two conditions should hold. (1) *A* originated as a "reproduction" (to give one example, as a copy, or a copy of a copy) of some prior item or items that, due in part to possession of the properties reproduced, have actually performed *F* in the past, and *A* exists because (causally historically because) of this or these prior performances. (2) *A* originated as the product of some prior device that, given its

As this function can only be satisfied if the action is performed under the right kind of epistemic circumstances, these circumstances are Normal conditions for its performance. Moreover, since it seems overwhelmingly plausible (and indeed implicit in STRATEGIC SUCCESS) that our epistemic modal judgments correlate with these epistemic circumstances, all requirements for biosemantic representation seem to be satisfied. Epistemic modal judgments represent relations to epistemic circumstances consisting of relevant statistical facts. (The problem of identifying relevant reference classes for the statistics that constitutes the epistemic circumstances is solved by the reproductive explanation implicated in STRATEGIC SUCCESS, whichever it is, as such an explanation will have to identify relevant reference classes. So we should be able to set Millikan's reference class worry aside.)

Applying IMPLEMENTATION FUNCTION to the likelihood judgments of our dehydrated gazelle, we should recognize that these judgments not only guide actions aimed at concrete goals, like getting water while avoiding predators. They also have as their function to implement an epistemic strategy, such as that of performing the action alternative that is ranked best when possible risks and benefits of each alternative are given weight in proportion to how strongly the gazelle's evidence suggest that they would accompany that alternative. The specific function of the gazelle's judgments of degrees of likelihood is to determine what weight the various possible outcomes are given. For this to contribute to successful strategy implementation in a Normal way, these degrees will have to correspond to the strengths of evidence and the statistical properties constituting these. Because the gazelle's judgments of likelihoods tend to correlate with these statistical properties, these judgments

circumstances, had performance of F as a proper function and that, under those circumstances, normally causes F to be performed by means of producing an item like A . (Millikan 1993: 13–14)

As defined, derived functions do not require that individual items of the sort produced by the mechanism from which they derive their function explain survival *individually*. Moreover, it would be a bad idea to so restrict derived functions. Presumably, reproduction often happens because of cumulative and often overdetermined effects of many such items.

count as representing whatever these statistical properties are. Analogous stories can be told about other epistemic modal judgments.

Given STRATEGIC SUCCESS, it thus seems that SUCCESS and biosemantics are compatible with and indeed support the idea that epistemic modal judgments represent some relevant epistemic circumstances.

What would this account of the representational content of epistemic modal *judgments* tell us about utterances expressing these judgments—about epistemic modal *claims*? Elsewhere I have suggested that such epistemic modal claims standardly have as their function to elicit corresponding epistemic modal judgments of hearers who face the same epistemic problem but whose epistemic access to the matter at hand is improved by that claim: “It might be that *p*” and “It is quite likely that *p*” have as their functions to elicit the judgment that *it might be that p* and the judgment that *it is quite likely that p*, respectively. (Björnsson and Almér 2010; Björnsson and Finlay 2010; Björnsson 2015; cf. e.g. Montminy 2012; Willer 2013; Lennertz 2014; Millikan this issue; Stojnić forthcoming). If this is correct, it invites an account of the representational content of epistemic modal claims analogous to that of epistemic modal judgments.

To get at such an account, we start by asking what this elicitation of epistemic modal judgments has done such that speakers have continued to make epistemic modal claims and addressees continued to form epistemic modal judgments in response. Two benefits stand out. The first is that in expressing the speaker’s epistemic modal judgment, the claim provides significant information about how the target possibility—what the speaker says might, must, or is likely to be the case—stands to the speaker’s evidence, thus letting the addressee form an epistemic modal judgment on a “speaker-expanded” evidential base. When it is clear that the speaker has better epistemic access to the target possibility than the addressee, that base is not only expanded, but also significantly improved. The other benefit is that even if the claim did not meaningfully improve the evidential base, it might have alerted the addressee to a possibility not previously considered. In particular, this is a common point of telling someone that something might be the case. In light of this, I suggest that:

DISCOURSE FUNCTION: Epistemic modal discourse has been reproduced because epistemic modal claims have elicited corresponding epistemic modal judgments and resulting actions that are supported by the addressees' speaker-expanded evidential base.

If DISCOURSE FUNCTION is correct, then it is a Normal condition for the function of the judgments and actions resulting from epistemic modal claims targeting p that the speaker's judgment adequately reflects the relevant relation between p and the addressees' speaker-expanded evidential bases. Suppose, for example, that the judgment that it might be that p represents that p is compatible with the evidence. Then the claim that it might be that p would represent that p is compatible with addressees' speaker-expanded evidential bases. Likewise, suppose that the judgment that it must be that p represents that p is ensured by the evidence, and the judgment that p is likely to a certain degree represents that p is supported to that degree by the evidence. Then the claim that it must be that p represents that p is ensured by addressees' speaker-expanded evidence, and the claim that p is likely to some degree that it has support to a corresponding degree from that evidence.⁶

⁶ On this account, might-claims characteristically have a "communal" content, that is, one involving the evidence of at least one judge other than the speaker. Following von Stechow and Gillies 2011, I have previously argued that because speakers often lack full justification for asserting things about the evidence of the hearer, communal readings are implausible (Björnsson and Almér 2010: 7–13). I now think that this was a mistake. Exactly because addressees tend to have privileged access to some constituents of the represented fact, it makes sense for this particular communicative practice to distribute responsibility for getting things right between speakers and hearers. I now take the justificatory norms of the practice to demand of speakers that (i) they've taken their evidence properly into account and (ii) do not obviously get the part of the content that concerns addressee's evidence wrong, while leaving it to hearers to more definitely confirm or deny the latter part.

6. Problems of dual contents

Once we recognize implementation functions and corresponding representations we also need to recognize that SUCCESS does not cut down on admissible contents in the way previously assumed. Think of the beaver splash. Because the diving underwater that it triggers has brought about its own reproduction by bringing about predator avoidance, and because the presence of predators is a Normal condition for such avoidance, the splash represents the presence of predators: *Predators here now*. But the system that produces splashes and diving has also been reproduced for balancing the costs of not diving when there are predators around and the cost of diving in the absence of predators, giving beavers an evolutionarily stable responsiveness to evidence of predators. Based on this, it seems plausible that the system has as its function to produce implementations of a strategy of diving under certain epistemic circumstances and not others. Given biosemantics and SUCCESS, the splash might thus represent a certain level of risk of predators: *Significant risk of predators here now*. But can we really accept the result that the splash means both that there are predators around, and that there is a significant risk that there are?

Similar questions arise for epistemic modal judgments. According to the arguments of section 3, judgments to the effect that p might, must, or is likely to be the case all represent that p is the case, and according to the argument of section 5, they also represent that p stands in some relevant relation to the epistemic circumstances.⁷ Thus, I have suggested that the judgment that it might be that p represents both that p and that p is compatible with the evidence. By the argument of section 5, it would similarly seem to follow that an ordinary non-modal, unguarded belief that p not only represents that p , but also that p is sufficiently supported by the evidence, as there might be reproductive explanations of a strategy of acting on unguarded beliefs supported by certain degrees

⁷ To say that some item represents two distinct conditions is not to say that it represents their conjunction. The latter requires that the conjunction is a Normal condition for the performance of a function of actions guided by that item. Cf. n. 3.

or kinds of evidential support. Likewise for ordinary descriptive utterances in the declarative: such utterances might not only represent what they are naturally understood as asserting, but also that the asserted content has a certain degree of evidential support, as there might be a reproductive explanation of the strategy of forming beliefs in response to such utterances when these are sufficiently evidentially supported. But can we really accept that modal judgments, regular beliefs, and ordinary assertions all have these combinations of epistemic and non-epistemic contents?

The phenomenon to come to grips with is this. On the one hand, we have good reproductive explanations pointing to concrete effects (predator avoidance, rehydration) of actions guided by various utterances, beliefs, and judgments, and corresponding representational contents. Call these contents “concrete”. On the other, we have good reproductive explanations pointing to the strategy implementation function of these same actions, and corresponding contents. Call these contents “strategic”. (Contents summarized in Table 2 below. Descriptions of strategic contents should be understood as rough indications rather than fully-fledged analyses.)

Table 2. Concrete and strategic contents

| Representation | Content | |
|---|----------|---|
| | Concrete | Strategic |
| Unguarded belief that p | p | p is sufficiently evidentially supported |
| Assertion that p | p | p is sufficiently evidentially supported |
| Judgment that it must be that p | p | p is ensured by the evidence |
| Judgment that it might be that p | p | p is compatible with the evidence |
| Judgment that p is likely to degree d | p | p is evidentially supported to degree d |

Though these combinations of concrete and strategic contents—these dual contents—might at first seem problematic, there are good reasons to embrace their possibility. It is widely recognized that things can have more than one content, as witnessed by familiar distinctions in philosophy of language between what is asserted and conventionally or conversationally implicated. It is also clear why a naturalist theory of representation can leave room for multiple contents. What such a theory

should provide is the principled identification of a theoretically interesting property. Specifically, the identification should provide enlightening unification, when possible, of phenomena that we intuitively explain with reference to contentful states, letting us see new similarities and differences among these phenomena, and important relations between these phenomena and others. If explanations of such phenomena can crisscross, as in the cases just considered, so can contents.

Judging from the argument thus far, if dual contents pose a problem, it is not by suggesting that SUCCESS and biosemantics fail to provide theoretical unification of the phenomena they do cover. Indeed, it is the application of the core ideas of SUCCESS and its biosemantics interpretation that gave rise to the problem of EPISTEMICS LOST, and it is their further application based on implementation functions that yields dual contents. Rather, if there is a worry here, it is that these core ideas overgenerate, attributing contents that are just implausible, or fail to make distinctions that are plainly there. In particular, one might worry that biosemantics cannot account for the following two phenomena:

INTUITIVE EPISTEMICS: Though biosemantics with SUCCESS and IMPLEMENTATION FUNCTION imply that epistemic modal judgments and ordinary unguarded beliefs alike have concrete as well as strategic, epistemic contents, only epistemic modal judgments are intuitively concerned with epistemic affairs.

UNINTUITIVE FACTIVITY: Though biosemantics and SUCCESS imply that judgments that *p might be the case* represent that *p*—even though they portray these judgments as in this sense factive—it seems, first, that one can accept that *p might be the case* without accepting that *p is the case*, and, second, that one's judgment that *p might be the case* can be true or correct even if *p is not* in fact the case. (Likewise for judgments that *it is likely that p*.)

Here I will briefly outline explanations of both these phenomena.

Start with the contrast between epistemic modal judgments and ordinary unguarded beliefs. Though both categories have concrete as well as strategic contents, there are contrasts both with respect to the formation of the relevant states and with respect to deliberation based on these states.

First, the formation of an unguarded belief that *p* does not characteristically involve reflection on what evidence there is for or against *p*. In most cases of occurrent belief, we just unreflectively

take it that p , and even in most cases of belief formation, we just perceive, hear, or read that something is the case without attending to facts about how the resulting belief is supported: we only attend to such facts when the belief is called into question. By contrast, judgments that p must, might, or is likely to be the case characteristically involve attention to evidence for or against p , and we understand the distribution of confidence or acceptance between p and its alternatives as determined by such evidence.

Second, and relatedly, deliberation based on the unguarded belief that p characteristically involves no consideration of evidence for or against p or attention to competing alternatives to p . Actions are simply adjusted to achieve their concrete goals given p , with no adjustments made to achieve these goals should p not be the case. By contrast, deliberation based on epistemic modal judgments will be concerned with what relative weight to give alternative possibilities in the pursuit of relevant concrete goals, and this weight will characteristically be guided by ongoing attention to evidence for or against these alternatives. Deliberation based on the judgment that p might be the case or that p is likely to some degree characteristically involves adjustments or openness to adjustments of actions so as to achieve concrete goals given alternatives to p . Even deliberation based on the judgment that it must be that p characteristically involves consideration of one or more alternatives, though these alternatives are given no weight in attempts to achieve concrete goals because they are taken to be incompatible with the evidence.

Because both the formation and deliberative use of epistemic modal judgments characteristically relates to the evidence in ways the formation and use of unguarded beliefs do not, the distinction between the categories remains clear even given the existence of dual contents.

Turn next to UNINTUITIVE FACTIVITY. Given that judgments that it might be that p represent not only that p is compatible with the evidence but also that p , why does it seem that one can accept that it might be that p without accepting that p ? And why does it seem that the judgment that it might be that p can be true or correct even when p is *not* the case?

To get at the answer to the first question, notice that if one accepts that it might be that p one is per hypothesis in a state of acceptance representing that p . This, though, is clearly not the second of

the states of acceptance that we have in mind when we think that one can (i) accept that p might be the case without (ii) accepting that p . The state we have in mind is the (ordinary, unguarded) *belief* that p . But it should be no mystery why one can judge that it might be that p without believing that p . Though both states prompt actions selected for their propensity to achieve concrete goals given p , their functional role differs in other regards: they respond differently to evidence and they guide actions in different ways.

If there is a real question for the biosemanticist here, it is why we intuitively think of “accepting that p ” as involving *belief* rather than some other state representing that p , such as judging that it might be that p . Here I think that we can point to the language used in attributing contentful states. The state that we intuitively describe as “accepting (or judging, or taking it) that S ”, where “ S ” stands in for some string of words in declarative form, is the kind of state constituting acceptance of the claim that would be made by uttering S unembedded.⁸ Thus, the state of accepting that there are lions nearby is the kind of state that constitute acceptance of the claim, “there are lions nearby”. Moreover, the acceptance of claims can plausibly be understood in terms of communicative function: to accept a claim is to make the judgment that the claim has as its function to elicit in hearers. Since the claim, “there are lions nearby” has as its communicative function to bring about the belief that there are lions nearby, this belief is what constitutes acceptance of the claim, and thus what constitutes accepting that there are lions nearby. We thus have a straightforward explanation of the first part of UNINTUITIVE FACTIVITY.

What remains to be explained, then, is why we think that the judgment that p *might be the case* can be correct or true even if p is not the case, given that the judgment represents that p . As is well

⁸ Talk of “the kind of state” constituting acceptance of the relevant claim is intentionally vague, meant to capture the fact that we can attribute judgments using a language that the judge does not understand, or using referring expressions the referent of which he would not recognize as applying to the objects of judgment. Cf. Millikan 1984: 207–20.

known from the philosophy of language, attributions of truth to assertions do not require the satisfaction of every content communicated by the assertion, only the asserted content. In particular, we generally assess the truth of assertions independently of the satisfaction of their strategic content. In the case of might-judgments, we have the converse situation: it seems that assessments of truth follow the strategic rather than concrete content. Ideally, we would like a general explanation of these phenomena. Here I will briefly recount a proposal that I have partly defended elsewhere, and indicate how it deals with the cases at hand. The core of the proposal is that attributions of truth and falsehood are concerned with the satisfaction of “fundamental standards” for relevant judgments (Björnsson 2015: 178–81; modified to fit this context):

TRUE/FALSE: We take it to be true (false) that *S*, take someone’s claims that *S* to be true (false), take someone’s judgment that *S* to be true (false) to the extent that we take it to conform to (violate) the fundamental standard for judgments constituting acceptance of *S*.

FUNDAMENTAL STANDARD: *F* is a fundamental standard for a judgment if and only if something is an evidential reason for (against) that judgment because, and to the extent that, it is evidence that the judgment satisfies (violates) *F*.

Take the judgment that there are lions around (at a certain distance from the waterhole, at a specific time). We can assess such a judgment against some epistemic standard, thinking that it was the correct judgment in light of the evidence. But that would not be a fundamental standard for the judgment. We only take something to provide evidential reasons for that judgment to the extent that it is evidence that there are lions around, and only take it to provide evidential reasons against the judgment to the extent that it is evidence that there are no lions around. We thus seem to be treating

F1: Judge that there are lions around if and only if there are lions around.

as the fundamental standard for the judgment. Given TRUE/FALSE, the standard attribution of truth and falsehood to judgments follows: the judgment that there are lions around is true if and only if there are lions around, false otherwise.

Contrast this with the judgment that there *might* be lions around. We clearly do not take the following to be a fundamental standard for this judgment:

F2: Judge that there might be lions around if and only if there are lions around.

In particular, moderately strong evidence that there are no lions around does not constitute moderately strong evidential reason against the judgment that there might be lions around, and overall weak evidence that there are lions around can constitute very strong evidential reasons for that judgment. Instead, the fundamental standards that we operate with for might-judgments seem to be the evidence-relative:

F3: Judge that there might be lions around if and only if the presence of lions is compatible with the evidence.

Evidence that the presence of lions is (not) compatible with the evidence is evidential reason of corresponding strength for (against) the judgment that lions might be present. The upshot of TRUE/FALSE, then, is that the concrete, factive content of might-judgments does not constitute what we think of as their truth-conditions; the strategic, epistemic content does. We thus have an account of the second part of UNINTUITIVE FACTIVITY.

What we have seen, then, is that dual contents are not in general problematic, that the distinction between epistemic modal judgments and other judgments with strategic epistemic contents remains clear, and that for the cases where specific dual contents seem implausible, that apparent implausibility can be explained without rejecting dual contents. Assuming that we otherwise have good reason to accept biosemantics and contents based on implementation functions, dual contents do not at this stage seem to pose any serious problem.

7. Concluding remarks

To account for the representational content of epistemic modal judgments, I have suggested that the reproduction of mechanisms producing epistemic modal judgments and resulting actions might

be explained in terms of the epistemic strategy that they implement. In virtue of such explanations, actions guided by epistemic modal judgments have implementation functions. Since epistemic facts constitute Normal conditions for the performance of such functions, epistemic modal judgments can represent epistemic facts even given SUCCESS. And since epistemic modal claims have as their function to elicit corresponding judgments, they too can represent epistemic facts.

At this stage, the proposal is highly programmatic. I have briefly explained how the representations generated by biosemantics given the recognition of implementation functions are compatible with natural thoughts about what it is to accept such contents and about the truth-conditions of epistemic modal judgments and ordinary non-modal beliefs. But I have not seriously discussed the nature of the relevant reproductive explanations, and my characterizations of the relevant epistemic facts are little more than placeholders. A serious development of the idea should say more about the explanations, spelling out in detail what sort of actions are guided by the various epistemic judgments and claims, and what properties of these actions are part of the strategic explanation of reproductive success. But this is as it should be with naturalistic semantic. Though we tend to have a rough intuitive grasp of the aspects of reality to which are adjusting our actions guided by our representations, getting at more precise representational contents requires increasingly systematic theorizing. Here I hope to have made it plausible that something along the lines presented might be true and thus worthy of further exploration.

Acknowledgements

This paper benefited from comments by participants of the award symposium for Ruth Millikan's 2017 Rolf Schock Prize and from Ruth's generous written comments on an early draft.

Bibliography

- Björnsson, Gunnar 2015: 'Disagreement, Correctness, and the Evidence for Metaethical Absolutism'. In *Oxford Studies in Metaethics*. Shafer-Landau, Russ (ed) Oxford University Press pp. 160–87.
- Björnsson, Gunnar and Almér, Alexander 2010: 'The Pragmatics of Insensitive Assessments: Understanding the Relativity of Assessments of Judgments of Personal Taste, Epistemic Modals, and More'. In *Formal Semantics and Pragmatics. Discourse, Context and Models*. Manhattan, Ks: New Prairie Press, Pp. 1–45. Partee, B.H., Glanzberg, M., & Skilters, J. (ed).
- Björnsson, Gunnar and Finlay, Stephen 2010: 'Metaethical Contextualism Defended'. *Ethics*, 121, pp. 7–36.
- Blackburn, Simon 2005: 'Success Semantics'. In *Ramsey's Legacy*. Lillehammer, Hallvard and Mellor, D. H. (eds) Oxford University Press pp. 22–36.
- Brandom, Robert B. 1994: 'Unsuccessful Semantics'. *Analysis*, 54, pp. 175-78.
- Häggqvist, Sören 2013: 'Teleosemantics: Etiological Foundations'. *Philosophy Compass*, 8, pp. 73-83.
- Lennertz, Benjamin 2014: 'Taking 'Might'-Communication Seriously'. *Analytic Philosophy*, 55, pp. 176-98.
- Millikan, Ruth Garrett 1984: *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- Millikan, Ruth Garrett 1989: 'Biosemantics'. *Journal of Philosophy*, pp. 86 281-97.
- Millikan, Ruth Garrett 1993: *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT-Pr.
- Millikan, Ruth Garrett 2004: *Varieties of Meaning*. A Bradford Book / MIT Press.
- Millikan, Ruth Garrett 2017: *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford University Press.
- Millikan, Ruth Garrett *this issue*: 'Biosemantics and Words That Don't Represent'. *Theoria*.
- Montminy, Martin 2012: 'Epistemic Modals and Indirect Weak Suggestives'. *Dialectica*, 66, pp. 583-606.
- Nanay, Bence 2013: 'Success Semantics: The Sequel'. *Philosophical Studies*, 165, pp. 151-65.
- Papineau, David 1993: *Philosophical Naturalism*. Oxford: Blackwell.
- Ramsey, F. P. 1927: 'Facts and Propositions'. *Proceedings of the Aristotelian Society*, 7, pp. 153-70.

Ryder, Dan, Kingsbury, Justine and Williford, Kenneth 2012: *Millikan and Her Critics*. John Wiley & Sons.

Stojnić, Una forthcoming: 'Content in a Dynamic Context'. *Noûs*, pp. n/a-n/a.

Tang, Weng Hong 2014: 'Success Semantics and Partial Belief'. *Journal of Philosophical Research*, 39, pp. 17-22.

von Fintel, Kai and Gillies, Anthony S. 2011: 'Might' Made Right'. In *Epistemic Modality*. Egan, Andy and Weatherson, Brian (eds) Oxford University Press pp. 108–30.

Whyte, J. T. 1990: 'Success Semantics'. *Analysis*, 50, pp. 149-57.

Willer, Malte 2013: 'Dynamics of Epistemic Modality'. *Philosophical Review*, 122, pp. 45-92.